




INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 11/00		A1	(11) International Publication Number: WO 99/63440
(21) International Application Number: PCT/US99/12000		(43) International Publication Date: 9 December 1999 (09.12.99)	
(22) International Filing Date: 2 June 1999 (02.06.99)			
(30) Priority Data: 60/087,733 2 June 1998 (02.06.98) US 09/140,174 25 August 1998 (25.08.98) US			(81) Designated States: AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GE, GH, GM, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).
(71) Applicant: ALLIEDSIGNAL INC. [US/US]; 101 Columbia Road, P.O. Box 2245, Morristown, NJ 07962-2245 (US).			
(72) Inventors: ZHOU, Jeffrey, Xiaofeng, 3908 Paul Mill Road, Ellicott City, MD 21042 (US). RODEN, Thomas, Gilbert, III, 3980 Hooper Road, New Windsor, MD 21776 (US). BOLDUC, Louis, P., 5359-E Columbia Road, Columbia, MD 21044 (US). PENG, Dar-Tzen, 6923 Newberry Drive, Columbia, MD 21045 (US). ERNST, James, W., 9501 Good Lion Road, Columbia, MD 21044 (US). YOUNIS, Mohamed, 5029 Columbia Road, Columbia, MD 21044 (US).			
(74) Agents: CRISS, Roger, H. et al.; AlliedSignal Inc., Law Dept. (Amy Olinger), 101 Columbia Road, P.O. Box 2245, Morristown, NJ 07962-2245 (US).			

(54) Title: METHOD AND APPARATUS FOR MANAGING REDUNDANT COMPUTER-BASED SYSTEMS FOR FAULT TOLERANT COMPUTING

(57) Abstract

A stand alone redundancy management system (RMS)(12) provides a cost-effective solution for managing redundant computer-based systems in order to achieve ultra-high system reliability, safety, fault tolerance and mission success rate. The RMS includes a cross channel data link (CCDL) module (24A) and a fault tolerant executive (FTE) module (13). The CCDL module provides data communication for all channels, while the FTE module performs system functions such as synchronization, data voting, fault and error detection, isolation and recovery. System fault tolerance is achieved by detecting and masking erroneous data through data voting, and system integrity is ensured by a dynamically reconfigurable architecture that is capable of excluding faulty nodes from the system and re-admitting healthy nodes back into the system.



FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		
EE	Estonia						

5

METHOD AND APPARATUS FOR MANAGING REDUNDANT COMPUTER-BASED SYSTEMS FOR FAULT TOLERANT COMPUTING**BACKGROUND OF THE INVENTION****1. Field of the Invention**

The present invention relates to computing environments, more particularly, it
10 relates to a method for managing redundant computer-based systems for fault-tolerant
computing.

2. Background of the Invention

Fault tolerant computing assures correct computing results in the existence of
faults and errors in a system. The use of redundancy is the primary method for fault
15 tolerance. There are many different ways of managing redundancy in hardware,
software, information and time. Due to various algorithms and implementation
approaches, most current systems use proprietary design for redundancy management,
and these designs are usually interwoven with application software and hardware. The
interweaving of the application with the redundancy management creates a more
20 complex system with significantly decreased flexibility.

SUMMARY OF THE INVENTION

It is therefore an object of the present invention to provide a method for managing
a redundant computer-based systems that is not interwoven with the application, and
provides additional flexibility in the distributed computing environment.

25

According to an embodiment of the present invention, the redundant computing
system is constructed by using multiple hardware computing nodes (or channels) and
installing a redundancy management system (RMS) in each individual node in a
distributed environment.

5 The RMS is a redundancy management methodology implemented through a set of algorithms, data structures, operation processes and design applied through processing units in each computing system. The RMS has wide application in many areas that require high systems dependability such as aerospace, critical control systems, telecommunications, computer networks, etc.

10 To implement the RMS, it is separated, physically or logically, from the application development. This reduces the overall design complexity of the system at hand. As such, the system developer can design applications independently, and rely on the RMS to provide redundancy management functions. The RMS and application integration is accomplished by a programmable bus interface protocol which connects the 15 RMS to application processors.

The RMS includes a Cross Channel Data Link (CCDL) module and a Fault Tolerant Executive (FTE) module. The CCDL module provides data communication for all channels while the FTE module performs system functions such as synchronization, voting, fault and error detection, isolation and recovery. System fault tolerance is 20 achieved by detecting and masking erroneous data through voting, and system integrity is ensured by a dynamically reconfigurable architecture that is capable of excluding faulty nodes from the system and re-admitting healthy nodes back into the system.

The RMS can be implemented in hardware, software, or a combination of both (i.e., hybrid), and works with a distributed system which has redundant computing 25 resources to handle component failures. The distributed system can have two to eight channels (or nodes) depending upon system reliability and fault tolerance requirements. A channel consists of a RMS and an application processor(s). Channels are interconnected together through the RMS's CCDL module to form a redundant system. Since individual applications within a channel do not have full knowledge of other 30 channel's activities, the RMS's provide system synchronization, maintain data consistency, and form a system-wide consensus of faults and errors occurring in various locations in the system.

5

BRIEF DESCRIPTION OF THE DRAWINGS

A more complete appreciation of this invention, and many of the attendant advantages thereof, will be readily apparent as the same becomes better understood by reference to the following detailed description when considered in conjunction with the accompanying drawings, in which like reference symbols indicate the same or similar 10 components, wherein:

FIG. 1 is a block diagram of the redundancy management system according to an embodiment of the present invention;

FIG. 2 is a block diagram of a three-channel RMS based fault tolerant system according to an exemplary embodiment of the present invention;

15 FIG. 3 is a state transition diagram of the redundancy management system according to an embodiment of the invention;

FIG. 4 is a block diagram of the redundancy management system, the application interaction and voting process according to an embodiment of the invention;

20 FIG. 5 is a schematic diagram of the context of the fault tolerant executive according to an embodiment of the present invention;

FIG. 6 is a block diagram of the voting and penalty assignment process performed by the fault tolerator according to an embodiment of the invention;

FIG. 7 is a schematic diagram of the context of the redundancy management system according to an embodiment of the invention;

25 FIG. 8 is a diagram of the cross channel data link message structure according to an embodiment of the invention;

FIG. 9 is a block diagram of the cross channel data link top level architecture according to an embodiment of the invention;

5 FIG. 10 is a block diagram of the cross channel data link transmitter according to an embodiment of the invention; and

FIG. 11 is a block diagram of the cross channel data link receiver according to an embodiment of the invention.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

10 According to an embodiment of the present invention, the redundancy management system (RMS) provides the following redundancy management functions: 1) Cross-channel data communication; 2) Frame-based system synchronization; 3) Data Voting; 4) Fault and error detection, isolation and recovery; and 5) a graceful degradation and self healing.

15 The cross-channel data communication function is provided by the CCDL module. The CCDL module has one transmitter and up to eight parallel receivers. It takes data from its local channel and broadcasts data to all channels including its own. Communication data is packaged into certain message formats and parity bits are used to detect transmission errors. All CCDL receivers use electrical-to-optical conversion in order to preserve electrical isolation among channels. Therefore, no single receiver failure can over drain current from other channel's receivers resulting in a common mode failure across the system.

20 The RMS is a frame-based synchronization system. Each RMS has its own clock and system synchronization is achieved by exchanging its local time with all channels and adjusting the local clock according to the voted clock. A distributed agreement algorithm is used to establish a global clock from failure by any type of faults, including Byzantine faults.

25 The RMS employs data voting as the primary mechanism for fault detection, isolation and recovery. If a channel generates data which is different from a voted majority, the voted data will be used as the output to mask the fault. The faulty channel

5 will be identified and penalized by a global penalty system. Data voting includes both application data and system status data. The RMS supports heterogeneous computing systems in which fault-free channels are not guaranteed to produce the exact same data (including data images) due to diversified hardware and software. A user specified tolerance range defines erroneous behavior should a data deviation occur in the voting
10 process.

The RMS supports a graceful degradation by excluding a failed channel from a group of synchronized, fault-free channels defining the operating set. A penalty system is designed to penalize erroneous behavior committed by any faulty channel. When a faulty channel exceeds its penalty threshold, other fault-free channels reconfigure
15 themselves into a new operating set that excludes the newly identified faulty channel. The excluded channel is not allowed to participate in data voting and its data is used only for monitoring purposes. The RMS also has the capability, through dynamic reconfiguration, to re-admit a healthy channel into the operating set. This self-healing feature allows the RMS to preserve system resources for an extended mission.

20 FIG. 1 shows a top-level block diagram of the RMS system according to an embodiment of the present invention. The RMS 12 includes a cross channel data link (CCDL) module 24a, and a fault tolerator executive module 13. The FTE 13 is resident on a VME card or other single board computer, and is connected to other cards in a system via the VME backplane bus or other suitable data bus. The RMS 12 is connected to other RMS's resident on each card via the CCDL module 24a. Each RMS includes its
25 own CCDL module for establishing a communication link between the respective RMS's. The establishment of a communication link via the CCDLs provides additional flexibility in monitoring the integrity of all cards in a system. By implementing a RMS on each computing node, and connecting the same to each other, system faults can be
30 detected, isolated, and dealt with more efficiently than other fault tolerant systems.

5 System Architecture

An exemplary three-channel RMS-base system architecture 10 according to an embodiment of the present invention is depicted in FIG. 2. In this architecture, the RMS interconnects three Vehicle Mission Computers (VMC) to form a redundant, fault tolerant system. Each VMC has a VME chassis with several single-board computers in it. The RMS 12 is installed in the first slot of VMC 1, and the communication between RMS and other application boards is through the VME backplane bus 14. Each VMC takes inputs from its external 1553 buses. The three main applications, Vehicle Subsystem Manager 16, Flight Manager 18 and Mission Manager 20, compute their functions and then store critical data in VME global memory (see FIG. 7) for voting.

15 Each RMS 12, 22 and 32, of the respective boards VMC1, VMC2 and VMC3, takes data via the VME bus and broadcasts the local data to other channels through the cross channel data link (CCDL) 24. After receiving three copies of data, the RMS will vote and write the voted data back to the VME global memory for use by the applications.

20 System Fault Tolerance

Each channel in the RMS is defined as a fault containment region (FCR) for fault detection, isolation, and recovery. Conventionally, an FCR usually has a territory bounded by natural hardware/software components. The key property of the FCR is its capability to prevent fault and error propagation into another region. Multiple faults occurring in the same region are viewed as a single fault because other regions can detect and correct the fault through the voting process. The number of simultaneous faults a system can tolerate depends upon the number of fault-free channels available in the system. For non-Byzantine faults, $N \geq 2f+1$ where N is the number of fault-free channels and f is the number of faults. If a system is required to be Byzantine safe, $N \geq 3f_B + 1$ where f_B is the number of Byzantine faults.

The RMS can tolerate faults with different time durations such as transient faults, intermittent faults and permanent faults. A transient fault has a very short duration and

5 occurs and disappears randomly. An intermittent fault occurs and disappears periodically with a certain frequency. A permanent fault remains in existence indefinitely if no corrective action is taken. In conventional fault tolerant systems design, rigorous pruning of faulty components can shorten fault latency, and thus, enhance the system's integrity. Nevertheless, immediate exclusion of transiently faulty components may decrease
10 systems resources too quickly, and jeopardize mission success. The fault tolerance of the RMS allows a user to program its penalty system in order to balance these two conflicting demands according to application requirements. Different penalties can be assigned against different data and system errors. High penalty weights for certain faults will result in rapid exclusion of faulty channels when such faults occur. Low penalty
15 weights against other faults will allow a faulty channel to stay in the system for a predetermined time so that it can correct its fault through voting.

According to the RMS system of the present invention, fault containment in three-node configuration excludes faulty channels when penalties exceed the user-defined exclusion threshold. A channel is re-admitted into the operating set when its
20 good behavior credits reach the re-admission threshold. Conflicts in application or channel data are resolved by mid-value selection voting.

In a two-node configuration, the RMS cannot detect or exclude a faulty node. As such, voting cannot be used to resolve conflicts. The application must determine who is at fault and take appropriate action.

25 **RMS Implementation**

As previously mentioned, the RMS has two subsystems, Fault Tolerant Executive (FTE) and Cross-Channel Data Link (CCDL). The FTE consists of five modules (FIG. 5): 1) a Synchronizer 80; 2) a Voter 58; 3) a Fault Tolerator (FLT) 84; 4) a Task Communicator (TSC) 46; and 5) Kernal (KRL) 52. The functions of these modules will
30 be described in the foregoing.

5 System Synchronization

The Synchronizer (SYN) 80 (FIG. 5) establishes and maintains channel synchronization for the system. It is required that, at any time, each individual RMS must be in, or operate in one of five states: 1)POWER_OFF; 2) START_UP; 3) COLD_START; 4) WARM_START; and 5) STEADY_STATE. FIG. 2 shows a state transition diagram of an individual RMS and its five states.

POWER_OFF (PF) is the state when the RMS is non-operational and the power source of the associated computer is off for any reason. When the RMS is powered-up, the RMS unconditionally transitions to START-UP.

START_UP (SU) is the state after the computer has just been powered up and when all system parameters are being initialized, RMS timing mechanisms are being initialized and the inter-channel communications links (i.e., CCDLs) are being established. When the start-up process is complete, the RMS unconditionally transitions to COLD_START.

COLD_START (CS) is the state in which the RMS cannot identify an existing Operating Set (OPS) and is trying to establish an OPS. The OPS is a group of nodes participating in normal system operation and voting. The RMS transitions from a WARM_START to COLD_START when less than two RMSs are in the OPS.

WARM_START (WS) is the state in which the RMS identifies the OPS containing at least 2 RMSs but the local RMS itself is not in the OPS.

25 STEADY_STATE (SS) is the state when the node of the RMS is synchronized with the OPS. A STEADY_STATE node can be in or out of the OPS. Each node in the OPS is performing its normal operation and voting. A node not included in the OPS is excluded from voting but its data is monitored by the OPS to determine its qualification for readmission.

30 In the Cold-Start, an Interactive Convergence Algorithm is used to synchronize channel clocks into a converged clock group which is the operating set (OPS). All

5 members are required to have a consistent view about their memberships in the OPS and they all switch to the Steady-State mode at the same time.

In the Steady-State mode, each channel broadcasts its local time to all channels through a System State (SS) message. Every channel dynamically adjusts its local clock to the global clock in order to maintain system synchronization. Since RMS is a frame-
10 synchronized system, it has a predetermined time window called the Soft-Error Window (SEW) that defines the maximum allowable synchronization skew. Each fault-free RMS should receive other SS messages in the time interval bounded by the SEW. Since RMS is used in a distributed environment, using a single SEW window has an inherent ambiguity in determining synchronization errors among participating channels. *See, P.*
15 *Thambidurai, A.M. Finn, R.M. Kieckhafer, and C.J. Walter, “Clock Synchronization in MAFT” Proc. IEEE 19th International Symposium on Fault-Tolerant Computing*, the entire content of which is incorporated herein by reference. To resolve the ambiguity, another time window known as a Hard-Error Window (HEW) is used. For example, if channel “A” receives channel “B’s” clock outside of “A’s” HEW, channel “A” reports a
20 synchronization error against channel “B”. However, if channel “B” sees that its own clock (after receiving its own SS message) is in the HEW, channel “B” reports that channel A has a wrong error report regarding synchronization. The ambiguity of mutually accusing channels needs to be resolved by other channel’s views about channel “B’s” clock. If channel “A” is correct, other channels should observe that channel “B’s”
25 clock has arrived at least outside of their SEW. Sustained by other channel’s error reports, the system can then identify channel “B” as the faulty channel. Otherwise, channel “A” is the faulty channel because of its deviation from majority view in the error report.

30 Warm-Start (WS) is half way between Cold-Start and Steady-State. A channel may be excluded from the OPS because of faults and errors. The excluded channel can go through reset and try to re-synchronize with the operating set in the Warm-Start mode. Once the channel detects that it has synchronized with the global clock of the operating set, it can switch to the Steady-State mode. Once in the Steady-State mode, the excluded channel is monitored for later re-admission into the OPS.

- 5 Time Synchronization within a VMC utilizes location monitor interrupts generated by RMS, and the VSM Scheduler uses frame boundary and Mid-frame signals for scheduling tasks.

Time Synchronization across the VMCs ensures source congruence. The CCDL time stamps RMS system data messages received over the 8Mbit data link. The FTE gets 10 the RMS system data from the VMCs and votes the time of these received messages, and adjusts CCDL local time to the voted value. The FTE then generates an interrupt on the synchronized frame boundary

System Voting

In RMS, voting is the primary technique used for fault detection, isolation, and 15 recovery. The RMS Voter (VTR) in the FTE votes on system states, error reports and application data. The voting of system states establishes a consistent view about system operation such as the membership in the OPS and synchronization mode. The voting on error reports formulate a consensus about which channel has erroneous behavior and what the penalty for these errors should be. The voting on application data provides 20 correct data output for the application to use. The data voting sequence is shown in FIG. 4.

The RMS data voting is a cyclic operation driven by a minor frame boundary. A minor frame is the period of the most frequently invoked task in the system. As 25 demonstrated in FIG. 4, a four channel system generates application data 40 in a minor frame and stores the data in a raw data shared memory 42 known as the application data table for RMS to vote. At the minor frame boundary 44, the Task Communicator (TSC) module 46 of the RMS uses the Data-Id Sequence Table (DST) 48 as pointers to read the data from application data table 42. The DST 48 is a data voting schedule which determines which data needs to be voted in each minor frame, and it also contains other 30 associated information necessary for voting. After reading the data, the TSC 46 packages the data into a certain format and sends the data to the CCDL 24. The CCDL broadcasts its local data to other channels while receiving data from other channels as well. When

- 5 the data transfer is completed, the Kernal (KRL) 52 takes the data from the CCDL 24 and stores the data in the Data Copies Table 56 where four copies of data are now ready for voting (i.e., 3 copies from other RMSs and one from the present RMS). The voter (VTR) 58 performs voting and deviance checks. A median value selection algorithm is used for integer and real number voting and a majority voting algorithm is used for binary and

10 discrete data voting. The data type and its associated deviation tolerance are also provided by the DST 48 which is used by the VTR 58 to choose a proper voting algorithm. The voted data 60 is stored in the Voted Data Table 62. At a proper time, the TSC module 46 reads the data from the voted table 62 and writes it back to the Application Data Table (or voted data shared memory) 66 as the voted outputs. Again,

15 the addresses of the output data are provided by the DST 48. For each voted data, a data conflict flag may be set in the Data Conflict Table 64 by the VTR 58 if the system has only two operating channels left and the VTR detects the existence of data disagreement. The Data Conflict Table 64 is located in a shared memory space so that the application software can access the table to determine if the voted data is valid or not.

20 Data Voting Options

Data Type	Description	Voting Algorithm	Voting Time Est.
Signed Integer	32 Bit Integer	Mid-Value Selection	6.0 secs
Float	IEEE single precision floating point	Mid-Value Selection	5.3 secs
Unsigned Integer	32 Bit word voted as a word (may be useful in voting status words)	Mid-Value Selection	6.0 secs
32 Bit Vector	32 bit word of packed booleans. Voted as 32 individual booleans	Majority Vote	12 secs

Table 1

Table 1 is an exemplary table of data voting options where the specified data types are IEEE standard data types for ANSI "C" language.

- 5 Fault Tolerator

By defining the Fault Containment Region as each channel, a FCR (i.e., channel) can manifest its errors only through message exchanges to other FCRs (channels). *See, J. Zhou, “Design Capture for System Dependability”, Proc. Complex Systems Engineering Synthesis and Assessment Workshop, NSWC, Silver Spring, MD, July 1992, pp107-119,* 10 incorporated herein by reference. Through voting and other error detection mechanisms, the fault tolerator (FLT) 84 (FIG. 5) summarizes errors into the 15 types shown in table 2. A 16 bit error vector is employed to log and report detected errors. The error vector is packaged in an error report message and broadcast to other channels for consensus and recovery action at every minor frame.

Error ID	Error Description	Detect ed By	Penalty Weight
E1	(Reserved)		
E2	A message is received with an invalid message type, node ID or data ID	CCDL	1 or TBD
E3	Horizontal or vertical parity error, incorrect message length, or message limit exceeded	CCDL	1 or TBD
E4	Too many Error Report or System State messages are received	CCDL	2 or TBD
E5	A non-SS message received within Hard-Error-Window	KRL	4 or TBD
E6	More than one of the same data has been received from a node	KRL	2 or TBD
E7	Missing SS message, or PRESYNC/SYNC messages do not arrive in the right order	SYN	2 or TBD
E8	An SS message does not arrive within the Hard-Error-Window (HEW)	SYN	4 or TBD
E9	An SS message does not arrive within the Soft-Error-Window (SEW)	SYN	2 or TBD
E10	An SS message was received with a minor and/or major frame number different from the local node	SYN	4 or TBD
E11	The CSS and/or NSS of the node do not agree with the Voted CSS and/or NSS	VTR	4 or TBD
E12	An error message has not been received from a node in this minor frame	VTR	2 or TBD
E13	Missing data message	VTR	2 or TBD
E14	The data value generated by a node is inconsistent with the voted value	VTR	2 or TBD
E15	The information contained in the error message from a node does not agree with that of the voted value	VTR	3 or TBD
E16	The number of errors accumulated for a node in one major frame has exceeded a preset limit	FLT	4 or TBD

15

Table 2 (Error Vector Table)

LEGEND:

CSS: Current System State indicated the nodes in the OPS in the current minor frame

- 5 NSS: Next System State indicated the nodes in the OPS in the next minor frame
OPS: Operating Set which is defined as the set of fault-free system nodes in Steady-State mode

TBD: To be Determined

CCDL: Cross Channel Data Link

10 KRL: Kernal

SYN: Synchronizer

VTR: Voter

FLT: Fault Tolerator

15 Referring to FIG. 6, the FLT 84 assesses a penalty 104 against a channel which is the source of errors. At every minor frame, all detected (reported) errors 100 are assigned with penalties using a penalty weight table 102, and the penalty sum is stored in a Incremental Penalty Count (IPC). The local IPC is assessed (104), and broadcast (106) to the other nodes via the CCDL. The FLT module votes on the IPC (108) and the voted

20 result is stored in a Base Penalty Count (BPC) 110. The IPC captures errors for a particular minor frame and the BPC captures cumulative errors for entire mission time.

After computing/storing the BPC (110), the IPC vector is cleared (112), and the BPC is broadcast (114) to the other nodes via the CCDL. The BPC is also voted (116) every minor frame and the FLT uses the voted BPC to determine whether a penalty assignment 25 and voting is required in order to ensure a consistent action among all fault-free channels for system reconfiguration. Once the voting on the BPC (116) is complete, the FLT determined whether a major frame boundary has been reached (118). If yes, the reconfiguration is determined (120). If the major frame boundary is not reached, the process returns to the error report 100, and continues from the beginning.

30 The system reconfiguration includes both faulty channel exclusion and healthy channel re-admission. If the Base Penalty Count (BPC) of a faulty channel exceeds a predetermined threshold, the RMS starts the system reconfiguration. During the reconfiguration, the system regroups the operating set to exclude the faulty channel.

Once a channel loses its membership in the operating set, its data and system status will 35 no longer be used in the voting process. The excluded channel needs to go through a reset process. If the reset process is successful, the channel can try to re-synchronize itself with the operating set and it can switch to the Steady-State mode if the synchronization is successful. An excluded channel can operate in the Steady-State

- 5 mode, but is still outside of the operating set. The channel now receives all system messages and application data from the nodes in the operating set.

10 All members in the operating set also receive messages from the excluded channel and monitor its behavior. The BPC of the excluded channel may be increased or decreased depending upon the behavior of the channel. If the excluded channel maintains fault-free operation, its BPC will be gradually decreased to below a predetermined threshold, and at the next major frame boundary, the system goes through another reconfiguration to re-admit the channel.

RMS and Application Interface

15 The current RMS implementation uses the VME bus and shared memory as the RMS and application interface. However, this is only one possible implementation and other communication protocol can also be employed to implement the interface. The main function of the TSC module 46 (FIG. 4) is to take data from designated communication media and package data into a certain format for the RMS to use. When a voting cycle is complete, the TSC takes the voted data and sends the data back to 20 application.

RMS Kernel

FIG. 5 shows a schematic diagram of the context of the fault tolerance executive (FTE) according to an embodiment of the invention. As shown, the Kernal 52 provides all of the supervisory operations for the RMS. The Kernal 52 manages the startup of 25 RMS, calling the appropriate functions to initialize the target processor as well as the loading of all initial data. During the startup process, the Kernal configures the CCDL module by loading the system configuration data and the proper operational parameters. The Kernal manages the transitions between the RMS operating nodes (i.e., Cold-Start, Warm-Start, and Steady- State) by monitoring the status of other RMS modules and 30 taking the appropriate actions at the correct times. The Kernal uses a deterministic scheduling algorithm such that all 'actions' are controlled by a self-contained time base.

- 5 At a given 'tick' in the time-base cycle, the predetermined actions for that tick are always executed. The Kernal 52 coordinates FTE functions based on the time tick. RMS activities, such as fault detection, isolation and recovery, are scheduled by the Kernal at the appropriate times in the RMS minor frame. If a RMS channel becomes faulty, the Kernal has the responsibility for restarting the channel at the proper time. All data
- 10 transfers between the RMS subsystems and between RMS and the application computer(s) are managed and scheduled by the Kernal. The Kernal directs the other modules to prepare various RMS messages and loads those messages into the CCDL for transmission at the Kernal's request. As messages are received by the CCDL, the Kernal extracts those messages and dispatches them to the correct module(s) for processing.
- 15 The Kernal runs in a loop, continuously executing each of the scheduled actions and monitoring the RMS status.

The Fault Tolerant Executive (FTE) provides Byzantine fault resilience for 4 or more nodes. Byzantine safe is provided for 3 nodes under the condition of source congruency. The FTE votes application data, removes/reinstates applications for FTE, 20 and synchronizes application and FTEs to < 100 sec skew.

In an exemplary embodiment, the FTE takes approximately 4.08 msec (40% utilization) to vote 150 words and perform operating system functions. The FTE memory is 0.4 Mbytes of Flash (5% utilization) and 0.6 Mbytes of SRAM (5% utilization). These values have been provided for exemplary purposes. It is to be 25 understood that one of ordinary skill in the art can alter these values without departing from the scope of the present invention.

RMS Context

FIG. 7 shows the RMS context or exchange structure between the RMS and VMC in the operating environment. The information being transferred within the VMC 30 includes the RMS System Data which is delivered at the RMS Frame Boundary, and includes information such as the minor frame number, the voted current/next system state for indicating who is operating in and out of the operating set, and a system conflict

- 5 flag for use in a two node configuration. The Data Conflict Table is used in a two node configuration for indicating an unresolvable data conflict on a peer data element basis. The Voted Output is the voted value for each data element submitted for voting from an operating set member. The RMS System Data, Data Conflict Table and Voted Output are transferred by the RMS to the Global shared memory that is in communication with
10 the local VMC in which the RMS is operating.

The Raw Output is data submitted to the RMS for voting by all nodes in Steady State mode. The Application Error Count is an optional capability of the system, and is transferred to the RMS for enabling an application to affect the error penalty assessed by the RMS in determining the operating set.

15 The Frame boundary information includes an interrupt to signal the beginning of a RMS frame. This signal frame synchronizes the FM, VSM, and MM. The Mid-Frame information is another interrupt which provides a signal 5 msecs from the beginning of the frame. The Application Data Ready information includes an interrupt generated by the RMS to signal the applications that voted data is waiting and can be retrieved and
20 processed. The System Reset is an optional control that the application can use on reset.

Cross Channel Data Link (CCDL)

The CCDL module provides data communication among channels. The data is packaged into messages and the message structure is shown in FIG. 8. As shown, the message structure includes a header, and various message types according to the types of messages being transmitted and received. Message type 0 is the structure of a data message; type 1 is the structure for a system state message; type 2 is the structure of a cold start message; and type 3 is the structure of an error report and penalty count
25 message.

Each CCDL has a transmitter and up to eight receivers. The CCDL top level
30 architecture, transmitter and receiver schematics are depicted in FIGS. 9-11. FIG. 9 shows a top level CCDL architecture with one transmitter 70, four receivers 72a-72d, and two interfaces 74a and 74b using a DY4 MaxPac mezzanine protocol. One interface 74b

- 5 facilitates data exchange between the base VME card and the CCDL memory, and the other interface 74a handles control logic and error reporting. When data needs to be transmitted, the CCDL interface 74b takes data from the base card and stores it into the 8bit transmitter memory 76. When data is received, the four receivers 72a-d process and store the received data in the four receiver memories 78a-d, respectively, one for each
10 node. The FTE then takes the data under the control of the CCDL. Since the CCDL is the only module which establishes physical connection among channels, it must enforce electrical isolation in order to guarantee Fault Containment Region for the system. The present CCDL uses the electrical-to-optical conversion to convert electrical signals to optical signals. Each receiver 72a-72d has a corresponding optical isolator 73a-73d to
15 provide the necessary isolation function. This enables every channel to have its own power supply, and all of them are electrically isolated from each other.

FIG. 10 shows a more detailed view of the transmitter 70 architecture in accordance with an embodiment of the present invention. When a "GO" command is issued by the FTE, the transmitter control logic 80 reads data from its 8 bit memory 76, forms the data into a 32 bit format, and appends a horizontal word to the end of the data.
20 The shift register circuit 82 converts the data into a serial bit string with vertical parity bits inserted into the string for transmission.

FIG. 11 illustrates how the serial data string is received from a transmitting mode and stored in its corresponding memory. The Bit Center logic 90 uses 6 system clock
25 (e.g., 48 MHZ) cycles to reliably log in one data bit. When the first bit of a data string is received, the Time Stamp logic 92 records the time for synchronization purposes. The shifter circuit 94 strips vertical parity bits and converts serial data into 8 bit format. An error will be reported should the vertical bit show transmission errors. The control logic 96 further strips horizontal parity from the data and stores it into the receiver memory
30 (e.g., 78a) according to the node number information attached with the data.

Both horizontal and vertical parity bits are attached to data messages in order to enhance communication reliability. Message format is verified by the CCDL and only valid messages are sent to the Kernal for further processing.

- 5

It should be understood that the present invention is not limited to the particular embodiment disclosed herein as the best mode contemplated for carrying out the present invention, but rather that the present invention is not limited to the specific embodiments described in this specification except as defined in the appended claims.

5 **WHAT IS CLAIMED IS:**

1. A method for managing redundancy computer-based systems having multiple hardware computing nodes (channels) comprising the steps of:
 - 10 providing a redundancy management system (RMS) in each computing node;
 - establishing a communication link between each RMS; and
 - implementing a fault tolerant executive (FTE) module in each RMS for managing faults and a plurality of system functions.
- 15 2. The method as claimed in claim 1, further comprising the step of synchronizing each computing node in the system, said step of synchronizing being performed by the FTE module and comprising the steps of:
 - 15 providing a clock in each RMS;
 - exchanging a local time in each RMS with all other nodes; and
 - adjusting the local clock of each respective RMS according to a voted system clock.
- 20 3. The method as claimed in claim 1, further comprising the steps of detecting faults/ errors in data generated in a node and preventing propagation of a detected fault/error in data generated in a node, said steps of detecting and preventing further comprising the steps of:
 - 20 voting on data generating by each node to determine whether data generated by one node is different from a majority; and
 - using the voted data as an output to mask a fault when data generated by a particular node is different from the voted majority.
- 30 4. The method as claimed in claim 1, wherein said step of providing a RMS in each computing node is performed independent of application development.

5

5. The method as claimed in claim 1, wherein said step of establishing is performed by incorporating a cross channel data link (CCDL) between the RMS of each computing node.

10 6. The method as claimed in claim 1, further comprising the steps of: defining each computing node (channel) as a fault containment region; detecting faults/errors in data generated in a computing node; and isolating a detected fault within the fault containment region to prevent propagation into another computing node.

15 7. The method as claimed in claim 6, wherein said step of detecting further comprises the step of voting on data generating by each node to determine whether data generated by one node is different from a voted majority.

20 8. The method as claimed in claim 7, wherein said step of isolating further comprises the step of using the voted data as an output to mask a fault when data generated by a particular node is different from the voted majority.

25 9. The method as claimed in claim 3, further comprising the steps of: identifying a faulty node in response to the result of data voting; penalizing the identified faulty node by a global penalty system; and excluding the identified faulty node from an operating set of nodes when the fault node's penalties exceed a user specified fault tolerance range.

30 10. The method as claimed in claim 9, further comprising the steps of: monitoring data on the excluded node to determine whether the excluded node qualifies for re-admission into an operating set; and

5 re-admitting the excluded node into the operating set when the monitoring indicates acceptable performance of the node within a predetermined threshold.

11. The method as claimed in claim 10, wherein the predetermined threshold is defined by a system operator.

10

12. A method for fault tolerant computing in computing environments having a plurality of computing nodes (channels), comprising the steps of:
implementing a redundancy management system (RMS) in each computing node independent from applications;
15 communicating between each RMS; and
 maintaining an operating set (OPS) of nodes for increasing fault tolerance of the computing environment.

20 13. The method as claimed in claim 12, wherein said step of communicating is performed on a cross channel data link (CCDL).

14. The method as claimed in claim 13, wherein said step of communicating further comprises the steps of:
25 interfacing the CCDL with the node of the respective RMS;
 providing a plurality of receivers in the CCDL for receiving data from each of the plurality of nodes, respectively;
 providing at least one transmitter in the CCDL for processing and transmitting the received data to a fault tolerant executive (FTE) resident in the RMS; and
30 providing at least one receiver memory and at least one transmitter memory for receiving and storing respective data as required.

5

15. The method as claimed in claim 12, wherein said step of maintaining an operating set of nodes is performed in a fault tolerant executive (FTE) resident in the RMS and further comprises the steps of:

10 receiving data from every node connected in the computing environment;
determining whether data received from any one node contains faults;
excluding a node which generated data that faulty with respect to other received data; and
re-configuring the operating set to not include the faulty node.

15 16. The method as claimed in claim 15, wherein said step of determining further comprises the steps of:

setting a tolerance range for faulty data;
voting on all received data from each node;
identifying a node having faulty data that exceeds the set tolerance range.

20 17. The method as claimed in claim 15, further comprising the steps of:
monitoring data on the excluded node; and
re-admitting the excluded node into the operating set when the monitored data indicates the correction of the faulty data on the excluded node.

25 18. The method as claimed in claim 16, wherein said step of voting is performed at every minor frame boundary in the data transmission.

30 19. The method as claimed in claim 15, wherein said step of re-configuring is performed at every major frame boundary in the data transmission.

5 20. An apparatus for managing redundancy computer-based systems
having multiple hardware computing nodes (channels) comprising:
 means for providing a redundancy management system (RMS) in each
 computing node;
 means for establishing a communication link between each RMS; and
10 means for implementing a fault tolerant executive (FTE) module in each
 RMS for managing faults and a plurality system functions.

15 21. The apparatus as claimed in claim 20, wherein said means for
establishing a communication link further comprises a cross channel data link
 connected to each redundancy management system in each computing node.

20 22. The apparatus as claimed in claim 20, further comprising:
 means for detecting faults/errors in data generated in any one node; and
 means for isolating a detected fault/error within the node from which the
 fault/error was generated.

23. The apparatus as claimed in claim 22, wherein said means for detecting
further comprises means for voting on data generated by each node for determining
whether data generated by one node is different from a voted majority.

24. The apparatus as claimed in claim 23, wherein said means for
25 isolating further comprises means for using the voted data to mask a fault generated
by one node that is different from the voted majority.

1/11





FIG. 1

2/11






FIG. 3





FIG. 4

5/11



5
FIG.

6/11




FIG. 6

7/11




FIG. 7

8/11

HEADER

MT	NID	MESSAGE COUNT
----	-----	---------------

MESSAGE TYPE 0: DATA MESSAGE

MT	NID	MESSAGE COUNT	RESERVED. (16)
		DATA ID	
		DATA VALUE	
		DATA ID	
		DATA VALUE	
		⋮	
		DATA ID	
		DATA VALUE	

MESSAGE TYPE 1: SYSTEM STATE MESSAGE

MT	NID	MESSAGE COUNT	FN. BITS	RESERVED. (12)
	ISW	NSS	CSS	RES. (8)
		PERIOD COUNTER		

MESSAGE TYPE 2: COLD START MESSAGE

MT	NID	MESSAGE COUNT	FN. BITS	RES. (4)	ISW
	ISW0	ISW1	ISW2		ISW3
	ISW4	ISW5	ISW6		ISW7

MESSAGE TYPE 3: ERROR REPORT AND PENALTY COUNT MESSAGE

MT	NID	MESSAGE COUNT	RESERVED. (16)
		ERROR VECTOR 0	IPC 0 BPC 0
		ERROR VECTOR 1	IPC 1 BPC 1
		ERROR VECTOR 2	IPC 2 BPC 2
		ERROR VECTOR 3	IPC 3 BPC 3
		ERROR VECTOR 4	IPC 4 BPC 4
		ERROR VECTOR 5	IPC 5 BPC 5
		ERROR VECTOR 6	IPC 6 BPC 6
		ERROR VECTOR 7	IPC 7 BPC 7
	AP/BIT 0	AP/BIT 1	AP/BIT 2 AP/BIT 3
	AP/BIT 4	AP/BIT 5	AP/BIT 6 AP/BIT 7

FIG. 8

9/11




FIG. 9

10/11




FIG. 10

11/11




FIG. 11

INTERNATIONAL SEARCH REPORT

International application No.

PCT/US99/12000

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) : G06F 11/00

US CL : 714/12

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 714/12, 10, 11

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

APS

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5,533,188 A (PALUMBO) 02 July 1996, col. 2, line 65 - col. 4, line 26, col. 5, line 25 - col. 6, line 61, col. 7, line 3 - col. 8, line 44.	1, 3-5, 12-15, 20-24
X	THAMBIDURAI, P. et al., "Clock Synchronization in MAFT," Proc. IEEE 19th Intl Symposium on Fault Tolerant Computing, 1989, pp. 142-149, page 142, col. 2, lines 15-30, page 143, col. 1, lines 13-24, page 148, col. 2, "Summary".	2
X	ZHOU, J., "Design Capture for System Dependability," Proc. Complex Systems Engineering Synthesis and Assessment Workshop, NSWC, Silver Spring, MD, July 1992, pp 107-119, page 108, third paragraph, page 115, second paragraph.	6-8

Further documents are listed in the continuation of Box C.

See patent family annex.

* Special categories of cited documents:	"T"	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
A document defining the general state of the art which is not considered to be of particular relevance	"X"	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
B earlier document published on or after the international filing date	"Y"	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*&*	document member of the same patent family
O document referring to an oral disclosure, use, exhibition or other means		
P document published prior to the international filing date but later than the priority date claimed		

Date of the actual completion of the international search

23 SEPTEMBER 1999

Date of mailing of the international search report

01 NOV 1999

Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231

Facsimile No. (703) 305-3230

Authorized officer

STEPHEN ELMORE

Telephone No. (703) 305-3847

Joni Hill